

Supporting Window Analytics over Large-scale Dynamic Graphs

Qi Fan[‡], Zhengkui Wang[§], Chee-Yong Chan[#], Kian-Lee Tan^{‡#}

[‡] *NUS Graduate School of Integrative Science and Engineering, National University of Singapore*

[#] *School of Computing, National University of Singapore*

[§] *Singapore Institute of Technology*

ABSTRACT

In relational DBMS, window functions have been widely used to facilitate data analytics. Surprisingly, while similar concepts have been employed for graph analytics, there has been no explicit notions of graph window analytic functions. In this paper, we formally introduce window queries for graph analytics. In such queries, for each vertex, the analysis is performed on a window of vertices defined based on the graph structure. In particular, we identify two instantiations, namely the *k-hop window* and the *topological window*. We develop two novel indices, *Dense Block index (DBIndex)* and *Inheritance index (I-Index)*, to facilitate efficient processing of these two types of windows respectively. Extensive experiments are conducted over both real and synthetic datasets with hundreds of millions of vertices and edges. Experimental results indicate that our proposed index-based query processing solutions achieve four orders of magnitude of query performance gain than the non-index algorithm and are superior over EAGR[16] wrt scalability and efficiency.

Categories and Subject Descriptors

H.2.4 [Systems]: Query processing; H.2, E.5 [Database]: Optimization

General Terms

Graph Database, Query Processing, Large Network

1. INTRODUCTION

Information networks such as social networks, biological networks and phone-call networks are typically modeled as graphs [8] where the vertices correspond to objects and the edges capture the relationships between these objects. For instance, in social networks, every user is represented by a vertex and the friendship between two users is reflected by an edge between the vertices. In addition, a user's profile can be maintained as the vertex's attributes. Such graphs contain a wealth of valuable information which can be analyzed to discover interesting patterns. For example, we can find the top-k influential users who can reach the most number of friends

within 2 hops. With increasingly larger network sizes, it is becoming significantly challenging to query, analyze and process these graph data. Therefore, there is an urgent need to develop effective and efficient mechanisms over graph data to draw out information from such data resources.

Traditionally, in relational DBMS, window functions have been commonly used for data analytics [7, 1]. Instead of performing analysis (e.g. ranking, aggregate) over the entire data set, a window function returns for each input tuple a value derived from applying the function over a window of neighboring tuples. For instance, users may be interested in finding each employee's salary ranking within the department. Here, each tuple's neighbors are essentially records from the same department.

Interestingly, the notion of window functions turns out to be not uncommon in graph data. For instance, in a social network, it is important to detect a person's social position and influence among his/her social community. The "social community" of the person is essentially his/her "window" comprising neighbors derived from his/her k-hop friends. However, as illustrated in this example, the structure of a graph plays a critical role in determining the neighboring data of a vertex. In fact, it is often useful to quantify a structural range to each vertex and then perform analytics over the range. Surprisingly, though such a concept of window functions has been widely used, the notion has not been explicitly formulated. In this paper, we are motivated to extend the window queries in traditional SQL for supporting graph analysis. However, the window definition in SQL is no longer applicable in a graph context, as it does not capture the graph structure information. Thus, we seek to formulate the notion of graph windows and to develop efficient algorithms to process them over large scaled graph structures.

We have identified two instantiations of graph windows, namely *k-hop* and *topological* windows. We first demonstrate these window semantics with the following examples.

EXAMPLE 1. (*k-hop window*) *In a social network (such as Linked-In and Facebook etc.), users are normally modeled as vertices and connectivity relationships are modeled as edges. In social network scenario, it is of great interest to summarize the most relevant connections to each user such as the neighbors within 2-hops. Some analytic queries such as summarizing the related connections' distribution among different companies, and computing age distribution of the related friends can be useful. In order to answer these queries, collecting data from every user's neighborhoods within 2-hop is necessary.*

EXAMPLE 2. (*Topological window*) *In biological networks (such as Argocyc, Ecocyc etc.[13]), genes, enzymes and proteins are vertices and their dependency in a pathway are edges. Because these networks are directed and acyclic, in order to study the pro-*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

tein regulating process, one may be interested to find out the statistics of molecules in each protein production pathway. For each protein, we can traverse the graph to find every other molecule that is in the upstream of its pathway. Then we can group and count the number of genes and enzymes among those molecules.

A common feature among these examples is that data aggregation is needed based on a set of vertices (which is the *graph window*) defined according to each vertex. To illustrate, in example 1, every user needs to gather data from its friends and friends-of-friends. The 2-hop neighbors form its window. Likewise, in example 2, every protein needs to count the number of particular type of genes preceding it in the regulating pathway. For every protein, the set of *preceding molecules* forms its window.

To support the analyses in the above-mentioned examples, we propose a new type of query, *Graph Window Query* (GWQ in short), over a data graph. GWQ is defined with respect to a graph structure and is important in a graph context. Unlike the traditional window in SQL, we identify two types of useful graph windows according to the graph structures, namely k-hop Window W_{kh} and Topological Window W_t . A k-hop window forms a window for one vertex by using its k-hop neighbors. k-hop neighbors are important to one vertex, as these are the vertices showing structural closeness as in Example 1. The k-hop neighbors window we define here is similar to the egocentric-network in network analysis [5] [16]. A topological window, on the other hand, forms a window for one vertex by using all its preceding vertices in a directed acyclic graph. The preceding vertices of one vertex are normally those which influence the vertex in a network as illustrated in Example 2.

To the best of our knowledge, existing graph databases or graph query languages do not directly support our proposed GWQ. There are two major challenges in processing GWQ. First, we need an efficient scheme to calculate the window of each vertex. Second, we need efficient solutions to process the aggregation over a large number of windows that may overlap. This offers opportunities to share the computation. However, it is non-trivial to address these two challenges.

For k-hop window like query, the state-of-the-art processing algorithm is EAGR [16]. EAGR builds an overlay graph including the shared components of different windows. This is done in multiple iterations, each of which performs the following. First, EAGR sorts all the vertices according to their k-hop neighbors based on their lexicographic order. Second, the sorted vertices are split into equal sized chunks each of which is further built as one frequent-pattern tree to mine the shared components. However, EAGR requires all the vertices' k-hop neighbors to be pre-computed and resided in memory during each sorting and mining operation; otherwise, EAGR incurs high computation overhead if the pre-computed structure needs to be shuffled to/from disk. This limits the efficiency and scalability of EAGR. For instance, a LiveJournal social network graph¹ (4.8M vertices, 69M edges) generates over 100GB neighborhood information for k=2 in adjacency list representation. In addition, the overlay graph construction is not a one-time task, but periodically performed after a certain number of structural updates in order to maintain the overlay quality. The high memory consumption renders the scheme impractical when k and the graph size increases.

In this paper, we propose *Dense Block Index* (DBIndex) to process queries efficiently. Like EAGR, DBIndex seeks to exploit common components among different windows to salvage partial work done. However, unlike EAGR, we identify the window simi-

larity utilizing a hash-based clustering technique instead of sorting. This ensures efficient memory usage, as the window information of each vertex can be computed on-the-fly. On the basis of the clusters, we develop different optimizations to extract the shared components which result in an efficient index construction.

Moreover, we provide another *Inheritance Index* (I-Index) tailored to topological window query. I-Index differentiates itself from DBIndex by integrating more descendant-ancestor relationships to reduce repetitive computations. This results in more efficient index construction and query processing.

Our contributions are summarized as follows:

- We propose a new type of graph analytic query, *Graph Window Query* and formally define two graph windows: k-hop window and topological window. We illustrate how these window queries would help users better query and understand the graphs under these different semantics.
- To support efficient query processing, we further propose two different types of indices: *Dense Block Index* (DBIndex) and *Inheritance Index* (I-Index). The DBIndex and I-Index are specially optimized to support k-hop window and topological window query processing. We develop the indices by integrating the window aggregation sharing techniques to salvage partial work done for efficient computation. In addition, we develop space and performance efficient techniques for index construction. Compared to EAGR [16], the state-of-the-art index method for k-hop window queries, our DBIndex is much more memory efficient and scalable towards handling the large-scale graphs.
- We perform extensive experiments over both real and synthetic datasets with hundreds of millions of vertices and edges on a single machine. Our experiments indicate that our proposed index-based algorithms outperform the naive non-index algorithm by up to four orders of magnitude. In addition, our experiments also show that DBIndex is superior over EAGR in terms of both scalability and efficiency. In particular, DBIndex saves up to 80% of indexing time as compared to EAGR, and performs well even when EAGR fails due to memory limitations.

2. RELATED WORK

Our proposed graph window functions (GWFs) for graph databases is inspired by the usefulness of window functions in relational analytic queries [23].

A window function in SQL typically specifies a set of partitioning attributes A and an aggregation function f . Its evaluation first partitions the input records based on A to compute f for each partition, and each input record is then associated with the aggregate value corresponding to the partition that contains the record. Several optimization techniques [7, 1] have also been developed to evaluate complex SQL queries involving multiple window functions.

However, the semantics and evaluation of window functions are very different between relational and graph contexts. Specifically, the partitions (i.e., subgraphs) associated with GWFs are not necessarily disjoint; thus, the evaluation techniques developed for relational context [7, 1] are not applicable to GWFs.

GWFs are also different from graph aggregation [24, 19, 8, 17] in graph OLAP. In graph OLAP, information in a graph are summarized by partitioning the graph's nodes/edges (based on some attribute values) and computing aggregate values for each partition. GWFs, on the other hand, compute aggregate values for each

¹Available at <http://snap.stanford.edu/data/index.html>, which is used in [16]

graph node w.r.t. the subgraph associated with the node. Indeed, such differences also arise in the relational OLAP context, where different techniques are developed to evaluate OLAP and window function queries.

In [20], the authors investigated the problem of finding the vertices that have top- k highest aggregate values over their h -hop neighbors. They proposed mechanisms to prune the computation by using two properties: First, the locality between vertices is used to propagate the upper-bound of aggregation; Second, the upper-bound value of aggregates can be estimated from the distribution of attribute values. However, all these pruning techniques are not applicable in our work, as we need to compute the aggregation value for every vertex. In such a scenario, techniques in [20] degrade to the non-indexed approach as described in Section 4.

Indexing techniques have been proposed to efficiently determine whether an input pair of vertices is within a distance of k -hops (e.g. k -reach index [9]) or reachable (e.g. reachability index [22]). However, such techniques are not efficient for computing the k -hop window or topological window for a set of n vertices with a time complexity of $O(n^2)$.

[16] proposed an EAGR system, which uses the famous VNM heuristic and Frequent-Pattern Tree to find the shared component among each vertex's neighborhoods. It starts by building an overlay graph as a bipartite graph representing the vertex-neighbor mapping. Then it aims to find the bi-cliques in the overlay graph. Each bi-clique represents a set of vertices whose neighborhood aggregates can be shared. Once a bi-clique is found, it is inserted back to the overlay graph as a virtual node to remove redundant edges. EAGR find bicliques in iterations. During each iteration, it sorts each vertices in overlay graph by their neighborhood information. Then the sorted vertices are split into equal-sized chunks. For each chunk, it then builds a FP-Tree to mining the large bi-cliques. As the algorithm iterates, the overlay graph evolves to be less dense.

The main drawback of EAGR is its demands of high memory usage on overlay construction. It requires the neighborhood information to be pre-computed, which is used in the sorting phase of each iteration. In EAGR the neighborhood information is assumed to be stored in memory. However, the assumption does not scale well for computing higher hop windows (such as $k \geq 2$). For instance, a LiveJournal social network graph² (4.8M nodes, 69M edges) generates over 100GB mapping information for $k=2$ in adjacency list representation. If the neighborhood information is resided in disk, the performance of EAGR will largely reduced. Similarly, if the neighborhood information is computed on-the-fly, EAGR needs to perform the computation in each iteration, which largely increases indexing time.

We tackle this drawback by adapting a hash based approach that clusters each vertex based on its neighborhood similarity. During the hashing, the vertex's neighborhood information is computed on-the-fly. As compared to sorting based approach, we do not require vertex's neighborhood to be pre-reside in memory. In order to reduce the repetitive computation of vertex's neighborhood, we further propose an estimation based indexing construction algorithm that only require a vertex's small hop neighborhood to be computed during clustering. As our experiments show, our proposed methods can perform well even when EAGR algorithm fails when neighborhood information overwhelms system's memory. To further reduce the neighborhood access, we adapted a Dense Block heuristic process each vertex in one pass. Experiments shows that the performance of our heuristic is comparable to EAGR's, but with much shorter indexing time.

²Available at <http://snap.stanford.edu/data/index.html>, which is used [16]

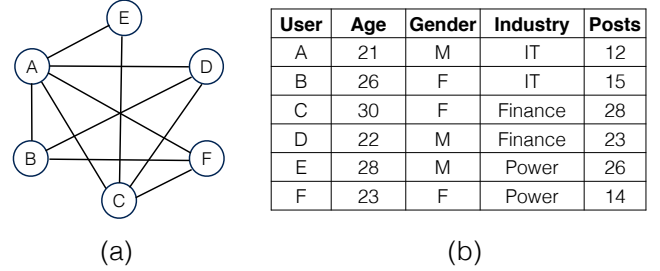


Fig. 1: Running Example of Social Graph. (a) provides the graph structure; (b) provides the attributes associated with the vertices of (a).

3. PROBLEM FORMULATION

In this section, we provide the formal definition of graph window query. We use $G = (V, E)$ to denote a directed/undirected data graph, where V is its vertex set and E is its edge set. Each node/edge is associated with a (possibly empty) set of attribute-value pairs.

Fig. 1 shows an undirected graph representing a social network that we will use as our running example in this paper. The table shows the values of the five attributes (User, Age, Gender, Industry, and Number of posts) associated with each vertex. For convenience, each node is labeled with its user attribute value; and there is one edge between a user X and another user Y if X and Y are connected in the social network.

Given a data graph $G = (V, E)$, a *Graph Window Function (GWF)* over G can be expressed as a quadruple (G, W, Σ, A) , where $W(v)$ denotes a *window specification* for a vertex $v \in V$ that determines the set of vertices in some subgraph of G , Σ denotes an *aggregation function*, and A denotes a *vertex attribute*. The evaluation of a GWF (G, W, Σ, A) on G computes for each vertex v in G , the aggregation Σ on the values of attribute A over all the vertices in $W(v)$, which we denote by $\Sigma_{v' \in W(v)} v'.A$.

Note that, in this paper, we focus on the attribute-based aggregation with distributive or algebraic aggregation functions. In other words, $W(v)$ refers to a set of vertices, and the aggregation function Σ operates on the values of attribute A over all the vertices in $W(v)$. Meanwhile, the aggregation function Σ is distributive or algebraic (e.g., sum, count, average), as these aggregation functions are widely used in practice.

In the following, we introduce two useful types of window specification (i.e., W), namely, *k-hop window* and *topological window*.

DEFINITION 1 (K-HOP WINDOW). *Given a vertex v in a data graph G , the k -hop window of v , denoted by $W_{kh}(v)$ (or $W(v)$ when there is no ambiguity), is the set of neighbors of v in G which can be reached within k hops. For an undirected graph G , a vertex u is in $W_{kh}(v)$ iff there is a α -hop path between u and v where $\alpha \leq k$. For a directed graph G , a vertex u is in $W_{kh}(v)$ iff there is a α -hop directed path from v to u ³ where $\alpha \leq k$.*

Intuitively, a k -hop window selects the neighboring vertices of a vertex within a k -hop distance. These neighboring vertices typically represent the most important vertices to a vertex with regard to their structural relationship in a graph. Thus, k -hop windows provide meaningful specifications for many applications, such as customer behavior analysis [3, 11], digital marketing [14] etc.

³ Other variants of k -hop window for directed graphs are possible; e.g., a vertex u is in $W_{kh}(v)$ iff there is a α -hop directed path from u to v where $\alpha \leq k$.

As an example, in Fig. 1, the 1-hop window of vertex E is $\{A, C, E\}$ and the 2-hop window of vertex E is $\{A, B, C, D, E, F\}$.

DEFINITION 2 (TOPOLOGICAL WINDOW). *Given a vertex v in a DAG G , the topological window of v , denoted by $W_t(v)$, refers to the set of ancestor vertices of v in G ; i.e., a vertex u is in $W_t(v)$ iff there is directed path from u to v in G .*

There are many directed acyclic graphs (DAGs) in real-world applications (such as biological networks, citation networks and dependency networks) where topological windows represent meaningful relationships that are of interest. For example, in a citation network where (X, Y) is an edge iff paper X cites paper Y , the topological window of a paper represents the citation impact of that paper [15, 12, 6].

As an example, Fig. 2 shows a small example of a Pathway Graph from a biological network. The topological window of E $W_t(E)$ is $\{A, B, C, D, E\}$ and $W_t(H)$ is $\{A, B, D, H\}$.




Fig. 2: Running Example of Pathway DAG. (a) provides the DAG structure; (b) provides the attributes associated with the vertices of (a).

DEFINITION 3 (GRAPH WINDOW QUERY). *A graph window query on a data graph G is of the form $GWQ(G, W_1, \Sigma_1, A_1, \dots, W_m, \Sigma_m, A_m)$, where $m \geq 1$ and each quadruple (G, W_i, Σ_i, A_i) is a graph window function on G .*

In this paper, we focus on graph window queries with a single window function that is either a k-hop or topological window. The evaluation of complex graph window queries with multiple window functions can be naively processed as a sequence of window functions one after another. We leave the optimization of processing multiple window functions for further study in the future.

4. DENSE BLOCK INDEX

A straightforward approach to process a graph window query $Q = (G, W, \Sigma, A)$, where $G = (V, E)$, is to dynamically compute the window $W(v)$ for each vertex $v \in V$ and its aggregation $\Sigma_{v' \in W(v)} v'.A$ independently from other vertices. We refer to this approach as *Non-Indexed* method.

Given that many of the windows would share many common nodes (e.g., the k-hop windows of two adjacent nodes), such a simple approach would be very inefficient due to the lack of sharing of the aggregation computations.

To efficiently evaluate graph window queries, we propose an indexing technique named *dense block index (DBIndex)*, which is both space and query efficient. The main idea of DBIndex is to try to reduce the aggregation computation cost by identifying subsets of nodes that are shared by more than one window so that the

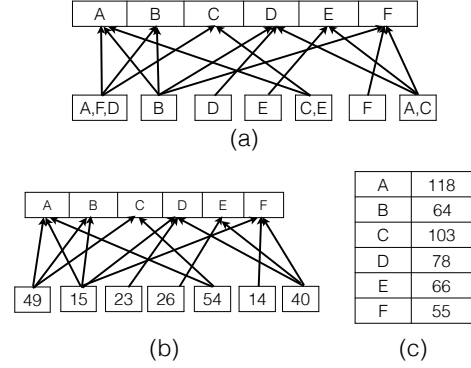


Fig. 3: Window Query Processing using DBIndex. (a) provides the DBIndex for 1-hop window query in Fig. 1; (b) shows the partial aggregate results based on the dense block; (c) provides the final aggregate value of each window.

aggregation for the shared nodes could be computed only once instead of multiple times.

For example, consider a graph window query on the social graph in Fig. 1 using the 1-hop window function. We have $W(B) = \{A, B, D, F\}$ and $W(C) = \{A, C, D, E, F\}$ sharing three common nodes A, D , and F . By identifying the set of common nodes $S = \{A, D, F\}$, its aggregation $\Sigma_{v \in S} v.A$ can be computed only once and then reuse to compute the aggregation for $\Sigma_{v \in W(B)} v.A$ and $\Sigma_{v \in W(C)} v.A$.

Given a window function W and a graph $G = (V, E)$, we refer to a non-empty subset $B \subseteq V$ as a *block*. Moreover, if B contains at least two nodes and B is contained by at least two different windows (i.e., there exists $v_1, v_2 \in V$, $v_1 \neq v_2$, $B \subseteq W(v_1)$, and $B \subseteq W(v_2)$), then B is referred to as a *dense block*. Thus, in the last example, $\{A, D, F\}$ is a dense block.

We say that a window $W(X)$ is *covered* by a collection of disjoint blocks $\{B_1, \dots, B_n\}$ if the set of nodes in the window $W(X)$ is equal to the union of all nodes in the collection of disjoint blocks; i.e., $W(X) = \bigcup_{i=1}^n B_i$ and $B_i \cap B_j = \emptyset$ if $i \neq j$.

To maximize the sharing of aggregation computations for a graph window query, the objective of DBIndex is to identify a small set of blocks \mathcal{B} such that for each $v \in V$, $W(v)$ is covered by a small subset of disjoint blocks in \mathcal{B} . Clearly, the cardinality of \mathcal{B} is minimized if \mathcal{B} contains a few large dense blocks.

Thus, given a window function W and a graph $G = (V, E)$, a DBIndex to evaluate W on G consists of three components in the form of a bipartite graph. The first component is a collection of nodes (i.e., V); the second component is a collection of blocks; i.e., $\mathcal{B} = \{B_1, \dots, B_n\}$ where each $B_i \subseteq V$; and the third component is a collection of links from blocks to nodes such that if a set of blocks $B(v) \subseteq \mathcal{B}$ is linked to a node $v \in V$, then $W(v)$ is covered by $B(v)$. Note that a DBIndex is independent of both the aggregation function (i.e., Σ) and the attribute to be aggregated (i.e., A).

Fig. 3(a) shows an example of a DBIndex wrt the social graph in Fig. 1 and the 1-hop window function. Note that the index consists of a total of seven blocks of which three of them are dense blocks.

4.1 Query Processing using DBIndex

Given a DBIndex wrt a graph G and a window function W , a graph window query $Q = (G, W, \Sigma, A)$ is processed by the following two steps. First, for each block B_i in the index, we compute the aggregation (denoted by T_i) over all the nodes in B_i ; i.e., $T_i = \Sigma_{v \in B_i} v.A$. Thus, each T_i is a partial aggregate value. Next,

for each window $W(v)$, $v \in V$, the aggregation for the window is computed by aggregating over all the partial aggregates associated with the blocks linked to $W(v)$; i.e., if $B(v)$ is the collection of blocks linked to $W(v)$, then the aggregation for $W(v)$ is given by $\sum_{B_i \in B(v)} T_i$.

Consider again the DBIndex shown in Fig. 3(a) defined wrt the social graph in Fig. 1 and the 1-hop window function. Fig. 3(b) shows how the index is used to evaluate the graph window query $(G, W, \text{sum}, \text{Posts})$ where each block is labeled with its partial aggregate value; and Fig. 3(c) shows the final query results.

4.2 DBIndex Construction

In this section, we discuss the construction of the DBIndex (wrt a graph $G = (V, E)$ and window function W) which has two key challenges.

The first challenge is the time complexity of the index construction. From our discussion of query processing using DBIndex, we note that the number of aggregation computations is determined by both the number of blocks as well as the number of links in the index; the former determines the number of partial aggregates to compute while the latter determines the number of aggregations of the partial aggregate values. Thus, to maximize the shared aggregation computations using DBIndex, both the number of blocks in the index as well as the number of blocks covering each window should be minimized. However, finding the optimal DBIndex to minimize this objective is NP-hard⁴. Therefore, efficient heuristics are needed to construct the DBIndex.

The second challenge is the space complexity of the index construction. In order to identify large dense blocks to optimize for query efficiency, a straightforward approach is to first derive the window $W(v)$ for each node $v \in V$ and then use this derived information to identify large dense blocks. However, this direct approach incurs a high space complexity of $O(|V|^2)$. Therefore, a more space-efficient approach is needed in order to scale to handle large graphs.

In this section, we present two heuristic approaches, namely *MC* and *EMC*, to construct DBIndex. The second approach EMC is designed to improve the efficiency of the first approach MC for constructing DBIndex (wrt k-hop window function) by using some approximation technique at the expense of possibly sacrificing the “quality” of the dense blocks (in terms of their sizes).

4.2.1 MinHash Clustering (MC)

To reduce both the time and space complexities for the index construction, instead of trying to identify large dense blocks among a large collection of windows, we first partition all the windows into a number of smaller clusters of similar windows and then identify large dense blocks for each of the smaller clusters. Intuitively, two windows are considered to be highly similar if they share a larger subset of nodes. We apply the well-known *MinHash based Clustering (MC)* algorithm [4] to partition the windows into clusters of similar windows. The MinHash clustering algorithm is based on using the *Jaccard Coefficient* to measure the similarity of two sets. Given the two window $W(v)$ and $W(u)$, $u, v \in V$, their *Jaccard Coefficient* is given by

$$J(u, v) = \frac{|W(u) \cap W(v)|}{|W(u) \cup W(v)|} \quad (1)$$

The *Jaccard Coefficient* ranges from 0 to 1, where a larger value means that the windows are more similar.

⁴ Note that a simpler variation of our optimization problem has been proven to be NP-hard [18].

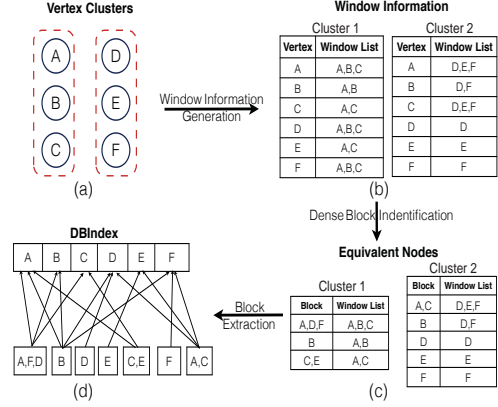


Fig. 4: DBIndex Construction over Social Graph in Fig. 1. (a) shows two clusters after MinHash clustering; (b) shows the window information of involved vertices within each cluster; (c) shows the dense blocks within each cluster; (d) provides the final DBIndex.

Our heuristic approach to construct DBIndex I operates as follows. Let $nodes(I)$, $blocks(I)$, and $links(I)$ denote, respectively, the collection of nodes, blocks, and links that form I . Initially, we have $nodes(I) = V$, $blocks(I) = \emptyset$, and $links(I) = \emptyset$.

Algorithm 1 CreateDBIndex

Require: Graph $G = (V, E)$, window function W
Ensure: DBIndex I

- 1: Initialize DBIndex I : $nodes(I) = V$, $blocks(I) = \emptyset$, $links(I) = \emptyset$
- 2: **for all** $v \in V$ **do**
- 3: Traverse G to determine $W(v)$
- 4: Compute the hash signature $H(v)$ for $W(v)$
- 5: **end for**
- 6: Partition V into clusters $\mathcal{C} = \{C_1, C_2, \dots\}$ based on hash signatures $H(v)$
- 7: **for all** $C_i \in \mathcal{C}$ **do**
- 8: **for all** $v \in C_i$ **do**
- 9: Traverse G to determine $W(v)$
- 10: **end for**
- 11: IdentifyDenseBlocks(I, G, W, C_i)
- 12: **end for**
- 13: **return** I

The first step is to partition the nodes in V into clusters using MinHash algorithm such that nodes with similar windows belong to the same cluster. For each node $v \in V$, we first derive its window $W(v)$ by an appropriate traversal of the graph G . Next, we compute a hash signature (denoted by $H(v)$) for $W(v)$ based on applying m hash functions on the set $W(v)$. Nodes with identical hash signatures are considered to have highly similar windows and are grouped into the same cluster. To ensure that our approach is scalable, we do not retain $W(v)$ in memory after its hash signature $H(v)$ has been computed and used to cluster v ; i.e., our approach does not materialize all the windows in the memory to avoid high space complexity. Let $\mathcal{C} = \{C_1, C_2, \dots\}$ denote the collection of clusters obtained from the first step, where each C_i is a subset of nodes.

The second step is to identify dense blocks from each of the clusters computed in the first step. The identification of dense blocks in each cluster C_i is based on the notion of node equivalence defined as follows. Two distinct nodes $u, v \in C_i$ are defined to be equivalent (denoted by $u \equiv v$) if u and v are both contained in the same set of windows; i.e., for every window $W(x)$, $x \in C_i$, $u \in W(x)$

iff $v \in W(x)$. Based on this notion of node equivalence, C_i is partitioned into blocks of equivalent nodes. To perform this partitioning, we need to again traverse the graph for each node $v \in C_i$ to determine its window $W(v)$ ⁵.

However, since C_i is now a smaller cluster of nodes, we can now materialize all the windows for the nodes in C_i in memory without exceeding the memory space. In the event that a cluster C_i is still too large for all its vertex windows to be materialized in main memory, we can further partition C_i into equal size sub-clusters. This re-partition process can be recursively performed until the sub-clusters created are small enough such that the windows for all nodes in the sub-clusters fit in memory.

Recall that a block B is a dense block if B contains at least two nodes and B is contained in at least two windows. Thus, we can classify the nodes in each C_i as either dense or non-dense nodes: a node $v \in C_i$ is classified as a *dense node* if v is contained in a dense block; otherwise, v is a non-dense node.

For each dense block B in C_i , we update the blocks and links in the DBIndex I as follows: we insert B into $blocks(I)$ if $B \notin blocks(I)$, and we insert (B, v) into $links(I)$ for each $v \in C_i$ where $B \subseteq W(v)$. If all the blocks in C_i are dense blocks, then we are done with identifying dense blocks in C_i ; otherwise, there are two cases to consider. For the first case, if all the nodes in C_i are non-dense nodes, then we also terminate the process of identifying dense blocks in C_i and update the blocks and links in the DBIndex I as before: we insert each non-dense block B into $blocks(I)$, and we insert (B, v) into $links(I)$ for each $v \in C_i$ where $B \subseteq W(v)$. For the second case, if C_i has a mixture of dense and non-dense nodes, we remove the dense nodes from C_i and recursively identify dense blocks in C_i following the above two-step procedure.

Note that since the blocks are identified independently from each cluster, it might be possible for the same block to be identified from different clusters. We avoid duplicating the same block in $blocks(I)$ by checking that a block B is not already in $blocks(I)$ before inserting it into $blocks(I)$. The details of the construction algorithm are shown as Algorithms 1, 2, and 3.

Algorithm 2 IdentifyDenseBlocks

Require: DBIndex I , Graph $G = (V, E)$, window function W , a cluster $C_i \subseteq V$

- 1: Partition C_i into blocks based on node equivalence
- 2: Initialize $DenseNodes = \emptyset$
- 3: **for all** dense block B **do**
- 4: Insert B into $blocks(I)$ if $B \notin blocks(I)$
- 5: Insert (B, v) into $links(I)$ for each $v \in C_i$ where $B \subseteq W(v)$
- 6: $DenseNodes = DenseNodes \cup B$
- 7: **end for**
- 8: **if** ($DenseNodes = \emptyset$) **then**
- 9: **for all** block B **do**
- 10: Insert B into $blocks(I)$ if $B \notin blocks(I)$
- 11: Insert (B, v) into $links(I)$ for each $v \in C_i$ where $B \subseteq W(v)$
- 12: **end for**
- 13: **else if** ($C_i - DenseNodes \neq \emptyset$) **then**
- 14: **if** ($C_i \neq DenseNodes$) **then**
- 15: RefineCluster($I, G, W, C_i - DenseNodes$)
- 16: **end if**
- 17: **end if**

Fig. 4 illustrates the construction of the DBIndex with respect

⁵ Note that although we could have avoided deriving $W(v)$ a second time if we had materialized all the derived windows the first time, our approach is designed to avoid the space complexity of materializing all the windows in memory at the cost of computing each $W(v)$ twice. We present an optimization in Section 4.2.2 to avoid the recomputation cost.

to the social graph in Fig. 1(a) and 1-hop window using the MC algorithm. First, the set of graph vertices are partitioned into clusters using MinHash clustering; Fig. 4(a) shows that the set of vertices $V = \{A, B, C, D, E, F\}$ are partitioned into two clusters $C_1 = \{A, B, C\}$ and $C_2 = \{D, E, F\}$.

For convenience, cluster 1 in Fig. 4(b) shows for each $v \in C_1$, the set of vertices whose windows contain v ; i.e., $\{u | v \in W(u)\}$. Similarly, Cluster 2 in Fig. 4 (b) shows for each $v \in C_2$, the set of vertices whose windows contain v . Consider the identification of dense blocks in cluster C_1 . As shown in Fig. 4 (c), based on the notion of equivalence nodes, cluster C_1 is partitioned into three blocks of equivalent nodes: $B_1 = \{A, D, F\}$, $B_2 = \{B\}$, and $B_3 = \{C, E\}$. Among these three blocks, only B_1 and B_3 are dense blocks. The MC algorithm then tries to repartition the window A, B, C using non-dense nodes in C_1 , (i.e., B_2). Since B_2 is the only non-dense node, it directly outputs. At the end of processing cluster C_1 , the DBIndex I is updated as follows: $blocks(I) = \{B_1, B_2, B_3\}$ and $links(I) = \{(B_1, \{A, B, C\}), (B_2, \{A, B\}), (B_3, \{A, C\})\}$. The identification of dense blocks in cluster C_2 is of similar process.

We find it is non-trivial to precisely analyze the complexity of Algorithm 1. Here, we only offer a brief analysis. Suppose the MinHash cost is H and the total cost for k-bounded BFS for all vertex is B , Lines 1-5 has the complexity of $O(H + B)$. Lines 7-10 has the complexity of $O(B)$. A single execution of Algorithm 2 has the complexity of $O(|V|)$, since we can simply partition nodes using hashing. Suppose the iteration runs for K times, the total cost for Algorithm 2 and Algorithm 3 is $O(K|V|)$. Therefore the overall complexity of Algorithm 1 is $O(H + 2 * B + O(K|V|))$. H depends on the number of vertex-window mappings for a given query and B depends on the graph structure and number of hops. As we demonstrate in Section 6, the H and B are the major contribution of the indexing time. To reduce the index time, we provide further optimization techniques.

Algorithm 3 RefineCluster

Require: DBIndex I , Graph $G = (V, E)$, window function W , a cluster $C \subseteq V$

- 1: **for all** $v \in C$ **do**
- 2: Compute the hash signature $H(v)$ for $W(v) \cap C$
- 3: **end for**
- 4: Partition C into clusters $\mathcal{C} = \{C_1, C_2, \dots\}$ based on hash signatures $H(v)$
- 5: **for all** $C_i \in \mathcal{C}$ **do**
- 6: IdentifyDenseBlocks(I, G, W, C_i)
- 7: **end for**

4.2.2 Estimated MinHash Clustering (EMC)

The MC approach described in the previous section requires the window of each node (i.e., $W(v), v \in V$) to be computed twice in order to avoid the high space complexity of materializing all the windows in main memory. For k-hop window function with a large value of k , the cost of graph traversals to compute the k-hop windows could incur a high computation overhead. Moreover, the cost of initial MinHash in MC approach equals to the initial number of vertex-window mappings, which is of the same order as graph traversal. For the larger hops, MinHash clustering would incur high computation cost.

To address these issues, we present an even more efficient approach, referred to as *Estimated MinHash Clustering (EMC)*, to optimize the construction of the DBIndex for k-hop window function with larger k .

The key idea behind *EMC* is based on the observation that for any two nodes $u, v \in V$, if their m -hop windows, $W_m(u)$ and $W_m(v)$, are highly similar and they are grouped into the same cluster, then it is likely that the n -hop windows of these two nodes, where $n > m$, would also be highly similar and grouped into the same cluster.

Using the above observation, we could reduce the overhead cost for constructing a DBIndex wrt a k -hop window function by clustering the nodes based on their k' -hop windows, where $k' < k$, instead of their k -hop windows.

To reduce the overhead of window computations, our *EMC* approach is similar to the MC approach except for the first round of window computations (line 3 in Algorithm 1): *EMC* uses lower hop windows to approximate k -hop windows for the purpose of clustering the nodes in V . Thus, the hash signatures used for partitioning V are based on lower hop windows. This approximation clearly has the advantage of improved time-efficiency as traversing and Minhashing on lower hop window is of order of magnitude faster. For the extreme case, adapting 1-hop window of a node v requires only accessing the adjacent nodes of v . The tradeoff for this improved efficiency is that the “quality” of the dense blocks might be reduced (in terms of their sizes). However, our experimental results show that this reduction in quality is actually only marginal which makes this approximation a worthy tradeoff.

4.2.3 Justification of Heuristic

In the following, we show the theoretical justification of our heuristic: the Jaccard coefficient is increasing wrt number of hops for a large class of graphs. We assume that the degree of vertices follows the same distribution, which is true in most real-networks and random graph models⁶. This implies we can analyze vertices with their neighborhoods structure using a unified way.

We use d_i to indicate the degree of vertex i . For any vertex pair (u, v) , their intersection on k -hop window consists of three part. We name them using $A = W_k(u) - W_k(v)$, $B = W_k(v) - W_k(u)$ and $C_k = W_k(u) \cap W_k(v)$. Clearly the Jaccard coefficient at k -hops can be expressed as follows:

$$J_k(u, v) = \frac{|C_k|}{|A_k| + |B_k| + |C_k|} \quad (2)$$

To deduct the relationship between $J_k(u, v)$ and $J_{k+1}(u, v)$, we prove the following lemma first:

THEOREM 4.1. *Let S be a collection of connected vertices, the number of newly discovered vertices by one-hop expansion from S is bounded by a function on $|S|$.*

PROOF. Consider a random variable Y_i indicate the newly discovered vertices from one-hop expansion from vertex i . Then the probability of $|Y_i| = y$ is can be analyzed as follows: there are d_i edges for vertex i . Since $|Y_i|$ is connected with S , one edge is fixed to link with a vertex in S . There are remaining $d_i - 1$ edges with y edges linked to the new vertices. In total, there are $\binom{|V|-1}{d_i-1}$ combinations with d_i edges. Therefore, the probability can be written as:

$$Prob(Y_i = y | v_i \in S) = \frac{\binom{|S|-1}{d_i-y-1} \binom{|V|-|S|}{y}}{\binom{|V|-1}{d_i-1}} \quad (3)$$

⁶E.g. Social network, Preferential Attachment model etc. follow power-law distribution. However, our analysis do not restrict on power-law distribution.

Thus, the expectation of Y_i is:

$$\begin{aligned} E(Y_i | v_i \in S) &= \sum (y * Prob(Y_i = y | v_i \in S)) \\ &= \sum_{y=1}^{d_i-1} \left(\frac{\binom{|S|-1}{d_i-y-1} \binom{|V|-|S|}{y}}{\binom{|V|-1}{d_i-1}} * y \right) \\ &= \sum_{y=1}^{d_i-1} \left(\frac{\binom{|S|-1}{d_i-y-1} \binom{|V|-|S|-1}{y-1}}{\binom{|V|-1}{d_i-1}} * (|V| - |S|) \right) \\ &= (|V| - |S|) * \sum_{y=1}^{d_i-1} \frac{\binom{|S|-1}{d_i-y-1} \binom{|V|-|S|-1}{y-1}}{\binom{|V|-1}{d_i-1}} \\ &= (|V| - |S|) * \frac{\binom{|V|-2}{d_i-2}}{\binom{|V|-1}{d_i-1}} = \frac{(|V| - |S|) * (d_i - 1)}{|V| - 1} \end{aligned} \quad (4)$$

Taking the expectation over all vertices in S , we can find the expectation of $E(Y|S) = \frac{(|V|-|S|)*(\bar{d}-1)}{|V|-1}$, where \bar{d} is the average degree of the graph. We then define the event $X = \cup_{i=1}^{|S|} Y_i$, i.e. X is the number of newly discovered vertices for one-hop expansion of entire S . By union bound, the expectation of X is:

$$\begin{aligned} E(X|S) &= E(\cup_{i=1}^{|S|} Y_i | S) \\ &\leq \sum_{i=1}^{|S|} E(Y_i | S) \\ &= \frac{|S|(|V| - |S|)(\bar{d} - 1)}{|V| - 1} = f(|S|) \end{aligned} \quad (5)$$

The bound is achieved when each vertices in S discovers non-overlapping neighbors, such as in the case of tree structure. Therefore, the newly discovered vertices are tightly bounded by a quadratic function on $|S|$. \square

We thus use $f(m)$ to denote the number of newly discovered vertices from a base set of m connected vertices. Since u, v have identical degree distribution, their expected value $S_u = E(|W_k(u)|)$ and $S_v = E(|W_k(v)|)$ is likely to be the same, i.e. $S_u \cong S_v$. We further use $\alpha = \frac{|C|}{|A|+|C|}$ to denote the portion of shared components in u 's k -hop neighborhood. Likewise, we use $\beta = \frac{|B|}{|B|+|C|}$ for $W_k(v)$. Since vertices have identical degree distribution, α and β are likely to be the same, i.e. $\alpha \cong \beta$. Now, the $J_{k+1}(u, v)$ for $(k+1)$ -hop can be represented as follows:

$$\begin{aligned} J_{k+1}(u, v) &= \frac{|C_{k+1}|}{|A_{k+1}| + |B_{k+1}| + |C_{k+1}|} \\ &= \frac{|C_k| + \alpha * f(S_u) + \beta * f(S_v)}{|A_k| + |B_k| + |C_k| + f(S_u) + f(S_v) - \Delta} \end{aligned} \quad (6)$$

, the Δ here is to compensate the doubly counted portion on the overlapping: $(A_{k+1} \cup B_{k+1}) \cap C_{k+1}$. Therefore, it subsequently follows:

$$\begin{aligned} J_{k+1}(u, v) &\geq \frac{|C_k| + \alpha * f(S_u) + \beta * f(S_v)}{|A_k| + |B_k| + |C_k| + f(S_u) + f(S_v)} \\ &\cong \frac{|C_k| + 2\alpha * f(S_u)}{|A_k| + |B_k| + |C_k| + 2 * f(S_u)} \end{aligned} \quad (7)$$




Fig. 5: I-Index Construction over the Pathway DAG in Fig. 2. (a) shows the DAG structure; (b) provides the inheritance relationship discovered during the index construction; (c) shows the final I-Index.

Due to the fact that $\alpha = \frac{|C_k|}{|C_k| + |A_k|} \geq \frac{|C_k|}{|A_k| + |B_k| + |C_k|}$, it follows:

$$\begin{aligned} \frac{\alpha f(S_u)}{f(S_u)} &\geq \frac{|C_k|}{|A_k| + |B_k| + |C_k|} \Leftrightarrow \\ \frac{|C_k| + 2\alpha * f(S_u)}{|A_k| + |B_k| + |C_k| + 2 * f(S_u)} &\geq \frac{|C_k|}{|A_k| + |B_k| + |C_k|} \\ &= J_k(u, v) \end{aligned} \quad (8)$$

Therefore, our analysis shows that $J_k(u, v)$ is most likely increasing for random graphs with identical degree distribution.

4.3 Handling Updates

In this section, we overview how our DBIndex is maintained when there are updates to the input graph. There are two types of updates for graph data: updates to the attribute values associated with the nodes/edges and updates to the graph structure (e.g., addition/removal of nodes/edges). Since the DBIndex is an index on the graph structure which is independent of the attribute values in the graph, the DBIndex is not affected by updates to the graph's attribute values.

The efficient maintenance of the DBIndex in the presence of structural updates is challenging as a single structural change (e.g., adding an edge) could affect many vertex windows. To balance the tradeoff between efficiency of index update and efficiency of query processing, we have adopted a two-phase approach to maintain the DBIndex. The first phase is designed to optimize update efficiency where the DBIndex is updated incrementally whenever there are structural updates to the graph. The incremental index update ensures that the updated index functions correctly but does not fully optimize the query efficiency of the updated index in terms of maximizing the shared computations. The second phase is designed to optimize query efficiency where the DBIndex is periodically reorganized to maximize share computations.

As an example of how the DBIndex is updated incrementally, consider a structural change where a new edge is added to the input graph. Let S denote the subset of graph vertices whose windows have expanded (with additional vertices) as a result of the insertion of the new edge. Let $W'(v)$ denote the set of additional vertices in the vertex window of v for each vertex $v \in S$. Based on the identified changes to the vertex windows (i.e., S and $\{W'(v) | v \in S\}$), we construct a secondary DBIndex which is then merged into the primary DBIndex. As the identified changes are small relative to the entire graph and collection of vertex windows, the construction and merging of the secondary index can be processed efficiently relative to an index reorganization to fully optimize query efficiency.

5. INHERITANCE INDEX

DBIndex is a general index that can support both k-hop as well as topological window queries. However, the evaluation of a topological window function, W_t , can be further optimized due to its containment feature. In other words, the window of a descendant vertex completely covers that of one of its ancestors. This feature can be formally formulated in the following theorem.

THEOREM 5.1. *In a DAG, if vertex u is the ancestor of vertex v , the topological window of v , $W_t(v)$ completely contains the window of u , $W_t(u)$, i.e., $W_t(u) \subset W_t(v)$.*

PROOF. In a DAG, if u is the ancestor of v , then $u \rightsquigarrow v$. $\forall w \in W_t(u)$, then $w \rightsquigarrow u$. As $u \rightsquigarrow v$, then $w \rightsquigarrow v$. Thus, $w \in W_t(v)$ and the theorem is proved. \square

Let us consider the BioPathway graph in Fig. 2 as an example. Fig. 5 (a) shows its abstract DAG. In (a), D is the ancestor of E. In addition, we can see that the window of D, $W_t(D)$ is $\{A, B, D\}$ and the window of E, $W_t(E)$ is $\{A, B, D, C, E\}$. It is easy to see that $W_t(D) \subset W_t(E)$.

Now, Theorem 5.1 provides us with opportunities for optimizing the space and computation of topological window queries. First, since the set of vertices corresponding to the window of a node, say u , is a superset of the set of vertices of its parent node, say v , there is no need to maintain the full set of vertices of the window at u . Instead, we only need to maintain the difference between $W_t(u)$ and $W_t(v)$. We note that in a DAG, it is possible for u to have multiple parents, v_1, \dots, v_k . In this case, the parent which has the smallest difference with u can be used; where there is a tie, it is arbitrarily broken. We refer to this parent as the *closest* parent. For instance, in Fig. 5 (a), instead of maintaining $\{A, B, D, C\}$ for $W_t(E)$, it can simply maintain the difference to $W_t(D)$ which is $\{C\}$. This is clearly more space efficient.

Second, using a similar logic, the aggregate computation at a node u can actually reuse the aggregate result of its closest parent, v . Referring to our example, the aggregation result of $W_t(D)$ can be simply passed or inherited to $W_t(E)$ and further aggregated with the difference set ($\{C\}$) in $W_t(E)$ to generate the aggregate value for $W_t(E)$. Fig. 5 (b) indicates the inheritance relationship that the values of the father can be inherited to the child in the tree.

Thus, we propose a new structure, called the **inheritance index**, *I-Index*, to support efficient processing of topological window queries. In *I-Index*, each vertex v maintains two information.

- The first information is the ID of the closest parent (say u) of v . We denote this as $PID(v)$.
- The second information is the difference between $W_t(v)$ and $W_t(u)$. We denote this as $WD(v)$.

With $PID(v)$, we can retrieve $W_t(u)$, and combining with $WD(v)$, we can derive $W_t(v)$. Likewise, we can retrieve the aggregation result of u which can be reused to compute v 's aggregation result.

Fig. 5(c) shows the I-index of our example in Fig. 5(a). In the figure, I-Index is represented in a table format; the second column is the PID and the third indicates the WD.

5.1 Index Construction

Building an I-Index for a DAG can be done efficiently. This is because the containment relationship can be easily discovered using a topological scan. Algorithm 4 lists the pseudo code for index creation. The scheme iterates through all the vertices in a topological order. For vertex v , the processing involves two steps. In the first step, we determine the closest parent of v . This is done by

comparing the cardinalities of the windows of v 's parents, and find the parent with largest value. The corresponding PID is recorded in the PID field of I -Index (Lines 7-12). In the second step, the window of v , $W_t(v)$, is pushed to its children (Lines 16-18). When the processing of v finishes, its window can be discarded. This frees up the memory space, which makes the scheme memory efficient.

We note that the complexity of Algorithm 4 is non-trivial to analyze. This is due to the difficulty of analyzing the number of ancestors of each vertex. Suppose the average number of ancestors for each vertex is H , then Algorithm 4 is of complexity $O(H|V| * d)$, where d is the average degree of the graph. This complexity is close to the output complexity. That is to gather the all vertex-window mapping, at least $O(H|V|)$ elements needs to be outputted. Thus the indexing time complexity is reasonably efficient.

We further note that the size of I -Index is hard to be precisely evaluated. This is due to the difficulty of analyzing the window difference. Assume the average size of window difference is D , then the size of I -Index is $O(D|V|)$. Although D can be as large $O(|V|)$, our experimental results indicate that the index size is always comparable to the graph size. We defer this discussion to section 6. Furthermore, it is possible to reduce the index size (should it be a concern) by employing compression techniques.

Algorithm 4 CreateI-Index

Require: Input graph: G

Ensure: Inheritance Index: $IIndex$

```

1:  $IIndex \leftarrow ()$ 
2:  $p \leftarrow ()$  {stores the window for each vertex}
3:  $c \leftarrow ()$  {stores the cardinality of window for each vertex}
4: for all  $v \in$  topological order do
5:    $WD \leftarrow -\infty$  {the window difference}
6:    $bestu \leftarrow nil$ 
7:   for all  $u \in v.parent$  do
8:     if  $c[u] > diff$  then
9:        $diff \leftarrow c[u]$ 
10:       $bestu \leftarrow u$ 
11:     end if
12:   end for
13:    $IIndex[v].WD \leftarrow WD$ 
14:    $IIndex[v].PID \leftarrow bestu$ 
15:    $p[v] \leftarrow p[v] \cup v$ 
16:   for all  $u \in v.child$  do
17:      $p[u] \leftarrow p[u] \cup p[v]$ 
18:   end for
19:    $c[v] \leftarrow |p[v]|$  {update window cardinality}
20:    $p[v] \leftarrow ()$  {release memory}
21: end for
```

5.2 Query Processing using I-Index

By employing the I-Index, window aggregation can be processed efficiently for each vertex according to the topological order. Algorithm 5 provides the pseudo code for the query processing. Each vertex v 's window aggregation value can be calculated by using the following formula:

$$\Sigma(W_t(v)) = \Sigma(W_t(v.PID), \Sigma(v.WD)) \quad (9)$$

where Σ is the aggregate function. As the vertex is processed according to the topological order, $W_t(v.PID)$ would have already been calculated while processing v 's parent and thus can be directly used for v without any recomputation. In general, v 's window aggregation is achieved by utilizing its parent's aggregate value and window difference sets. This avoids repeated aggregate computation

and achieves the goal of computation sharing between a vertex and its parent. In so doing, the computation overhead can be further reduced. Take the index provided in Fig. 5 (c) as an example, assume the query wants to calculate the sum value over each window for every vertex. As a comparison, the number of add operations are 33, 22, 16 for the cases without any index, with DBIndex and with I-Index index respectively.

Algorithm 5 QueryProcessingOverIIndex

Require: Input graph G , aggregate function Σ , inheritance index $IIndex$

Ensure: w {The aggregation result of each vertex}

```

1:  $w \leftarrow ()$ 
2: for all  $v \in$  topological order do
3:    $u \leftarrow IIndex[v].PID$ 
4:    $WD \leftarrow IIndex[u].WD$ 
5:    $S \leftarrow v.val$ 
6:    $S \leftarrow \Sigma(S, w[u])$ 
7:   for all  $t \in WD$  do
8:      $S \leftarrow \Sigma(S, t.val)$ 
9:   end for
10:   $w[v] \leftarrow S$ 
11: end for
12: return  $w$ 
```

As the query processing in Algorithm 5 basically scans the I -Index, the query complexity essentially correlates to the index size. As we shown in the experiment session, the query can be performed efficiently in various graph conditions. We defer the discussion to Section 6.

5.3 Handling Updates

We cover the updates handling in this section. For attribute updates, I -Index is not affected since I -Index is only structure related. Structure updates on I -Index consists of node updates and edge updates. It is easy to handle the case where an isolated node is added or delete, since it does not affect any other nodes in the graph. Adding (resp. deleting) a node with edges can be done via edge insertions (resp. deletions). Here we focus on describing single edge insertion and deletion. We use I to denote the I -index and $I(v)$ to denote the index entry of v .

During an update of edge $e(s, t)$, there are two types of vertices are affected. The first type contains single node t , which is the endpoint of edge e . The second type of nodes contains all the descendants of t . There are altogether four special cases needs to be consider during the updates. We illustrate the four cases with aids of Fig. 6. In Fig. 6, the dashed edge $e(s, t)$ is the edge to be updated (added or deleted). The node u is the *lowest common ancestor* (LCA) of t and s . The cloud shape A and B are nodes in between of u, s and u, t . Since u is the LCA, A and B are thus disjoint. Bold edge $e(c, d)$ indicate that $I(d).PID$ is c . We distinguish and handle the four cases as follows:

Case I ($I(t).PID = s$): As shown in Fig. 6 (a), during deletion, t needs to choose a parent from A to be its $I(t).PID$. $I(t).WD$ needs to be updated accordingly. In this case, no insertion needs to be considered since if there is no edge between s and t , $I(t).PID$ cannot be s .

Case II ($I(t).PID \neq s$): As shown in Fig. 6 (b), during insertion, B needs to be excluded from $I(t).WD$. During deletion, any node in B that cannot reach t needs to be included in $I(t).WD$. Since A and B are disjoint, every node in B needs to be removed from $I(t).WD$

Case III ($t \rightsquigarrow I(v).PID$): As shown in Fig. 6 (c), during insertion, any node in B needs to be removed from $I(v).WD$. During deletion, any node in B that reaches v but cannot reach a needs to be added to $I(v).WD$.

Case IV ($t \nrightarrow I(v).PID$): As shown in Fig. 6 (d), during insertion, any node in B that cannot reach r needs to be included into $I(v).WD$. During deletion, any node in B that cannot reach v needs to be excluded from $I(v).WD$.

During structure updates, essential operations are computing A , B and performing reachability queries. Computing A , B can be done via existing techniques such as [2, 10] while reachability query can be supported by indexing methods such as [21]. We defer exploring for more efficient updating algorithms to future work

6. EXPERIMENTAL EVALUATION

In this section, we present a comprehensive experimental evaluation of our solutions using several real-world information networks and various synthetic datasets. Since the focus of this paper is on query processing efficiency, we do not evaluate the efficiency of index updates. All experiments are conducted on an Amazon EC2 r3.2xlarge machine⁷, with an 8-core 2.5GHz CPU, 60GB memory and 320GB hard drive running with 64-bit Ubuntu 12.04. As the source code of EAGR is not available, we implemented it and used it as a reference in our comparative study. All algorithms are implemented in Java and run under JRE 1.6.

Name	Type	# of Vertices	# of Edges
LiveJournal1	undirected	3,997,962	34,681,189
Pokec	directed	1,632,803	30,622,564
Orkut	undirected	3,072,441	117,185,083
DBLP	undirected	317,080	1,049,866
YouTube	undirected	1,134,890	2,987,624
Google	directed	875,713	5,105,039
Amazon	undirected	334,863	925,872
Stanford-web	directed	281,903	2,312,497

Table 1: Large Scale Real Data

Datasets. For real datasets, we use 8 information networks which are available at the Stanford SNAP website⁸: LiveJournal1, Pokec, Orkut, DBLP, YouTube, Google, Amazon and Stanford-web. The detail description of these datasets is provided in Table 1.

For synthetic datasets, we use two widely used graph data generators. We use the *DAGGER* generator [21] to generate all the synthetic DAGs and the SNAP graph data generator at the Stanford SNAP website to generate non-DAG datasets. For each dataset, each vertex is associated with an integer attribute.

Query. In all the experiments, the window query is conducted by using the SUM() as the aggregate function over the integer attribute in each dataset.

6.1 Comparison between MC and EMC

We first compare the effectiveness of the two DBIndex construction algorithms: MinHash Clustering (MC) and Estimated MinHash Clustering (EMC). We look at the index construction time, index sizes and query performance. All these experiments are conducted based on two real world datasets: Amazon and Stanford-web. For both datasets, we run a series of k -hop queries.⁹ For queries

⁷<http://aws.amazon.com/ec2/pricing/>

⁸<http://snap.stanford.edu/snap/index.html>

⁹For the Stanford-web graph, which is directed, the k -hop windows are directed k -hop windows where $u \in W(k)$ if there is a directed

with hop count larger than 1, EMC uses 1-hop information for the initial clustering.

Index Construction. Figs. 7 (a) and (b) compare the index construction time between MC and EMC when we vary the windows from 1-hop to 4-hop for the Amazon and Stanford-web graphs respectively. To better understand the time difference, the construction time is split into two parts: the MinHash cost (EMC-hash or MC-hash) and the breadth-first-search traversal (to compute the k -hop window) cost (EMC-bfs or MC-bfs). The results show the same trend for the two datasets. We made several observations. First, as the number of hops increases, the indexing time increases as well. This is expected as a larger hop count results in a larger window size and the BFS and computation time increase correspondingly. Second, as the hop count increases, the difference between the index time of EMC and that of MC widens. For instance, as shown in Figs. 7(a) and (b), for the 4-hop window queries, compared to MC, EMC can save 62% and 66% construction time for the Amazon and Stanford-Web datasets respectively. EMC benefits from both the low MinHash cost and low BFS cost. From Figs. 7 (a) and (b), we can see that the MinHash cost of MC increases as the number of hops increases, while that for EMC remains almost the same as the 1-hop case. This shows that the cost of MinHash becomes more significant for larger windows. Thus, using 1-hop clustering for larger hop counts reduces the MinHash cost in EMC. Similarly, as EMC saves on BFS cost for k -hop queries where $k > 1$, the BFS cost of EMC is much smaller than that of MC as well.

Index Size. Figs. 7 (c) and (d) present the effect of hop counts on the index size for the Amazon and Stanford-web datasets respectively. The y-axis shows the index ratio which is the index size over the original graph size. The insights we derive are: First, the index size is rather small compared to the original graph - it varies from 3% to 12% of the original graph for the Amazon dataset and from 8% to 22% for the Stanford-web dataset. Second, the index size decreases as the number of hops increases. While this appears counter-intuitive initially, it is actually reasonable - a larger hop results in a bigger window, which leads to more dense blocks. Third, the index ratio of EMC is slightly larger than that of MC for larger hop count. This indicates that MC can find more dense blocks than EMC to reduce the index size. Fourth, the index ratio on the Amazon dataset is much smaller than the ones on the Stanford-web dataset. This is because the Amazon dataset is undirected while the Stanford-web dataset is directed. For the Stanford-web dataset, since we use directed k -hop windows, the window size is naturally smaller.

Query Performance. Figs. 8 (a) and (b) present the query time of MC and EMC on the two datasets respectively as we vary the number of hops from 1 to 4. To appreciate the benefits of an index-based scheme, we also implemented a *Non-indexed* algorithm which computes window aggregate by performing k -bounded breadth first search for each vertex individually in real time. In Figures 8 (a) and (b), the execution time shown on the y-axis is in log scale. The results show that the index-based schemes outperform the non-index approach by four orders of magnitude. For instance, for the 4-hop query over the Amazon graph, our algorithm is 13,000 times faster than the non-index approach. This confirms that it is necessary to have well-designed index support for efficient window query processing. By utilizing DBIndex, for these graphs with millions of edges, every aggregation query can be processed in just between 30ms to 100ms for the Amazon graph and between 60ms to 360ms for the Stanford-web graph. In addition, we can see that as the number of hops increases, the query time decreases. This is

path of at most k hops from vertex v to vertex u .

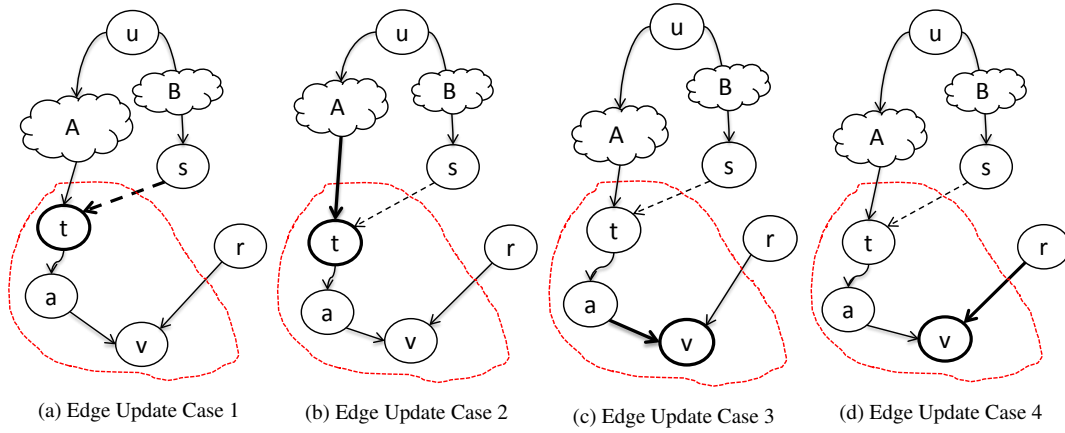


Fig. 6: Updates on *I*-Index. Cloud shape indicate the nodes in the subgraph between the endpoint nodes. The dashed circle indicate the affected range of updates. The bold arrow indicates the *PID* field *I*-Index

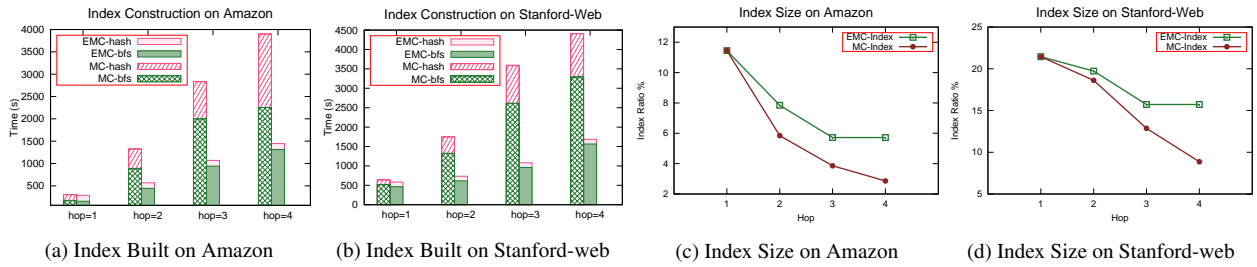


Fig. 7: Index Construction Analysis for EMC and MC. (a) and (b) depict the index time for the Amazon and Stanford-web networks; (c) and (d) shows the index size for the Amazon and Stanford-web datasets

the case because a larger hop count eventually results in a larger number of dense blocks where more (shared) computation can be salvaged. Furthermore, we can see that the query time of EMC is slightly longer than that of MC when the number of hops is large. This is expected as EMC does not cluster based on the complete window information; instead, it uses only partial information derived from the 1-hop windows. However, the performance difference is quite small even for 4-hop queries- for the Amazon dataset, the difference is only 20ms; and for the Stanford-web graph, the difference is 35ms. For small number of hops, the time difference is even smaller. This performance penalty is acceptable as tens of milliseconds time difference will not affect user's experience. As EMC is significantly more efficient than MC in index construction, EMC may still be a promising solution to many applications. As such, in the following sections, we adopt EMC for DBIndex in our experimental evaluations.

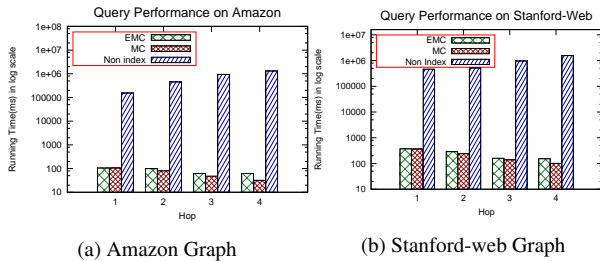


Fig. 8: Query Performance Comparison of MC and EMC

6.2 Comparison between DBIndex and EAGR

In this set of experiments, we compare DBIndex and EAGR [16]

using both large-scale real and synthetic datasets. Like [16], for each dataset, EAGR is run for 10 iterations in the index construction.

6.2.1 Real Datasets

We first study the index construction and query time performance of DBIndex and EAGR for 1-hop and 2-hop windows using 6 real datasets: DBLP, Youtube, Livejournal, Google, Pokec and Orkut. The results for 1-hop window and 2-hop window are presented in Figs. 9 and 10 respectively. As shown in Figs. 9(a) and 10(a), both DBIndex and EAGR can build the index for all the real datasets for 1-hop but EAGR ran out of the memory for 2-hop window queries on LiveJournal and Orkut datasets. This further confirms that EAGR incurs high memory usage as it needs to build the FPT and maintain the vertex-window mapping information. We also observe that DBIndex is significantly faster than EAGR in index creation. We emphasize that the time is shown in logarithmic scale. For instance, for Orkut dataset, EAGR takes 4 hours to build the index while DBIndex only takes 33 minutes.

Fig. 9 (b) and Fig. 10 (b) show the query performance for 1-hop and 2-hop queries respectively. The results indicate that the query performance is comparable. For most of the datasets, DBIndex is faster than EAGR. In some datasets (e.g. Orkut and Pokec), DBIndex performs 30% faster than the EAGR. We see that, for 1-hop queries on Youtube and LiveJournal datasets and 2-hop queries on Youtube dataset, DBIndex is slightly slower than EAGR. We observe that these datasets are very sparse graphs where the intersections among windows are naturally small. For very sparse graphs, both DBIndex and EAGR are unable to find much computation sharing. In this case, the performance of DBIndex and EAGR is very close. For instance, in the worse case for livejournal,

DBIndex is 9% slower than EAGR where the actual time difference remains tens of millionseconds.

Another insight we gain is that as expected, compared to Fig. 9 (b), 2-hop query runs faster for both algorithms. This is because there is more computation sharing for 2-hop window query.

In summary, DBIndex takes much shorter time to build but offers comparable, if not much faster, query performance than EAGR.




Fig. 9: Comparison between DBIndex and EAGR for 1-hop query

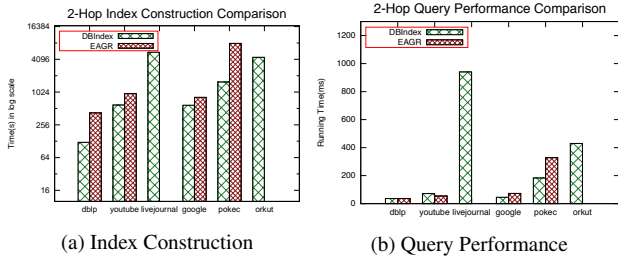


Fig. 10: Comparison between DBIndex and EAGR for 2-hop query


6.2.2 Synthetic Datasets

To study the scalability of DBIndex under large-scale networks, we generated synthetic datasets using the SNAP generator.

Impact of Number of Vertices. First, we study how the performance changes when we fix the degree¹⁰ at 10 and vary the number of vertices from 2M to 10M. Figs. 11 (a) and (b) show the execution time for index construction and query performance respectively. From the results, we can see that DBIndex outperforms EAGR in both index construction and query performance. For the graph with 10M vertices and 100M edges, the DBIndex query time is less than 450 milliseconds. Moreover, when the number of vertices changes from 2M to 10M, the query performance only increases 3 times. This shows that DBIndex is not only scalable, but offers acceptable performance..


Impact of Degree over Sparse Graphs. Our proposed DBIndex is effective when there is significant overlap between windows of neighboring nodes. As such, it is interesting to study how it performs for sparse graph where the nodes may not share many common neighbors. So, in these experiments, we study the impact of degree when the graph is relatively sparse. We fix the number vertices of 2M and vary the vertex degree from 5 to 30. Figs. 12 (a) and (c) present the results on index construction for 1-hop and 2-hop queries respectively. For 1-hop queries, as degree increases, the time for index construction also increases. However, the index creation time of DBIndex increases much slower than EAGR. This is because EAGR incurs relatively more overhead to handle multiple FPT creation and reconstruction. For 2-hop queries, EAGR failed to run. This is because even for a degree 5 sparse graph, the initial

¹⁰Degree means average degree of the graph. The generated graph is of Erdos-Renyi model




(a) Index Construction (b) Query Performance

Fig. 11: Impact of number of vertices



(a) Index Construction (b) Query Performance



(c) Index Construction (d) Query Performance

Fig. 12: Impact of Degree over Sparse Graphs over 2M vertices. (a) and (b) are the results for 1-hop query; (c) and (d) are the results for 2-hop query.

vertex-mapping can be as large as 90GB in a linked list manner, which exceeds the available memory. Note that the size becomes even larger when it is stored in a matrix manner. Therefore, we can only show the results of DBIndex. In Fig. 12 (c), indexing time of DBIndex increases as the degree increases. This is expected as a bigger degree increases the overhead of graph traversal time to collect the window.

Fig. 12 (b) and (d) show the results on query time for 1-hop and 2-hop queries respectively. We observe a similar pattern for the index construction time: for 1-hop queries, the query time increases with increasing degree but at a much slower rate than EAGR; for 2-hop queries, we observe in Fig. 12 (d) that the query performance of DBIndex hovers around 100ms, which is much smaller than that of 1-hop query performance. This is because there are more dense blocks in the 2-hop case, in which case the query time can be faster compared to the 1-hop case.

Impact of Degree over Dense Graphs. We study the impact of degree over very dense graphs with 200k vertices when the degree changes from 80 to 200. Figs. 13 (a) and (c) show the execution time for index construction for 1-hop and 2-hop queries respectively. From the results, we can see that DBIndex performs well for dense graphs as well. As the degree increases, EAGR's performance degrades much faster than DBIndex. For 2-hop queries, as shown in Figs. 13 (b) and (d), EAGR is only able to work on the dataset with degree 80 due to the memory issue. Even though the number of vertices is relatively small (only 200k), the number of edges is very large when the degree becomes big (e.g. 40M edges with degree of 200).

Figs. 13 (b) and (d) show the results on query performance for 1-hop and 2-hop queries respectively. The results are consistent with that for sparse graphs - DBIndex is superior over EAGR.

In summary, the insight we obtain is that the scalability of EAGR is highly limited by its approach to build the index over the vertex-window mapping information. EAGR is limited by two factors: the graph size and the number of hops. DBIndex achieves better scalability as it does not need to create a large amount of intermediate data in memory.

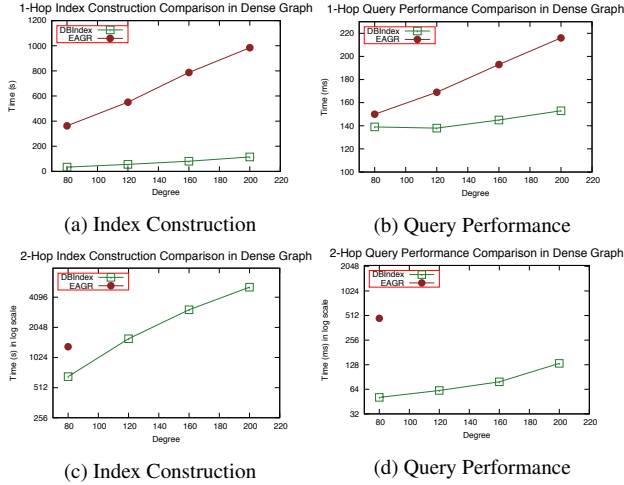


Fig. 13: Impact of Degree over Dense Graphs over 200K Vertices. (a) and (b) are the results for 1-hop query; (c) and (d) are the results for 2-hop query.

6.3 Evaluation of I-Index

In this set of experiments, we evaluate I-Index. All the datasets are generated from the DAGGER generator.

Impact of Degree. First, we evaluate the impact of degree changes when we fix the number of vertex as 30k and 60k. We compare DBIndex with I-Index. In the query results, we also implement one non-index algorithm which dynamically calculates the window and then performs the aggregation. For indexing time, as shown in Figures 14 (a) and (c), as the index size increases, both the indexing time of DBIndex and I-Index increase. However I-Index is more efficient than DBIndex, this is due to the special containment optimization used. We observe that the index construction time is almost the same as the one time non-index query time. In other words, we can use one query time to create the index which is able to provide much faster query processing for subsequent queries. In terms of query performance, shown in Figs. 14 (b) and (d), the non-index approach is, on average, 20 times slower than the index-based schemes. I-Index outperforms DBIndex by 20% to 30%. The results clearly show that I-Index outperforms DBIndex for topological window in both index construction and query performance. Therefore, in the following experiments, we only present the results for I-Index.

Impact of Number of Vertices. Next, we study how the performance of I-Index is affected when we fix the degree and vary the number of vertices from 50k to 350K. Figs. 15 (a) and (c) show the index construction time when we fix the degree to 10 and 20 respectively. From the results, we see that the construction time increases while the number of vertex increases and the construction time of a high degree graph is longer than that for low degree graphs. Figs. 15 (b) and (d) show the query time when we fix the degree to 10 and 20 respectively. As shown, the degree affects the query processing time - when the degree increases, the query time increases as well. We also observe that the query time is increasing linearly when the number of vertices increases. This shows the

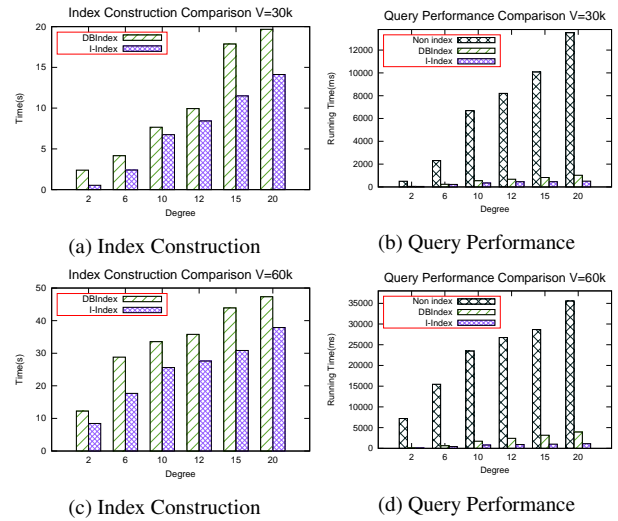


Fig. 14: Impact of Degree. (a) and (b) are the results for 30K vertices; (c) and (d) are the results for 60K vertices.

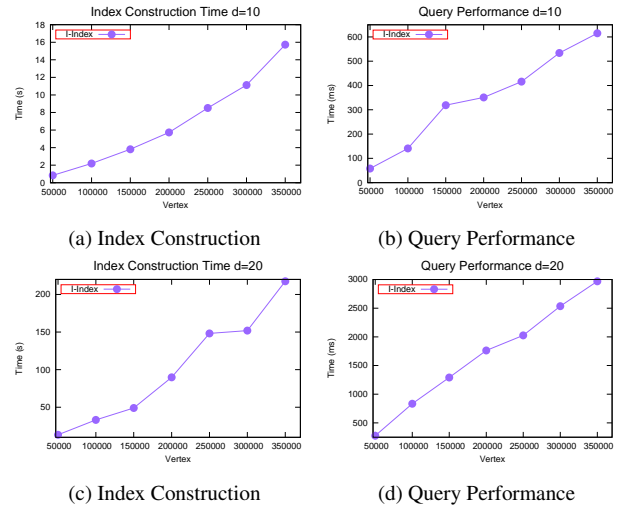


Fig. 15: Impact of the number of vertices with a fixed degree. (a) and (b) are the results for the graphs with degree 10; (c) and (d) are the graphs with degree 20.

I-Index has good scalability.

Index Size. Fig. 16 presents the index size ratio (i.e. size of index divided by the size of original graph) under different degrees from 3 to 20. There are four different sizes of data used with 100k, 150k, 200k and 300k vertices. For every vertex setting, the index size maintains the same trend in various degrees. The index size is linear to the input graph size. As a graph gets denser, the difference field of the I-Index effectively shrinks. Thus, the index size in turn becomes smaller, which explains the bends in the figure.

7. CONCLUSION AND FUTURE WORK

In this paper, we have proposed a new type of graph analytic query, *Graph Window Query*. We formally defined two instantiations of graph windows: k-hop window and topological window. We developed the Dense Block Index (DBIndex) to facilitate efficient processing of both types of graph windows. In addition, we also proposed the Inheritance Index (I-Index) that exploits a containment property of DAG to further improve the query perfor-




Fig. 16: Index Ratio of Inheritance-Index

mance of topological window queries. Both indices integrate window aggregation sharing techniques to salvage partial work done, which is both space and query efficient. We conducted extensive experimental evaluations over both large-scale real and synthetic datasets. The experimental results showed the efficiency and scalability of our proposed indices.

There remain many interesting research problems for graph window analytics. As part of our future work, we plan to explore structure-based window aggregations which are complex than attribute-based window aggregations. In structure-based aggregations, $W(v)$ refers to a subgraph of G instead of a set of vertices, and the aggregation function Σ (e.g., centrality, PageRank, and graph aggregation [19, 24]) operates on the structure of the subgraph $W(v)$.

8. ACKNOWLEDGMENT

Qi Fan is supported by NGS Scholarship. This work is supported by the MOE/NUS grant R-252-000-500-112 and AWS in Education Grant award.

9. REFERENCES

- [1] S. Bellamkonda, H.-G. Li, U. Jagtap, Y. Zhu, V. Liang, and T. Cruanes. Adaptive and big data scale parallel execution in oracle. In *VLDB'13*.
- [2] M. A. Bender, M. Farach-Colton, G. Pemmasani, S. Skiena, and P. Sumazin. Lowest common ancestors in trees and directed acyclic graphs. *Journal of Algorithms*, 2005.
- [3] E. J. Briscoe, D. S. Appling, R. L. Mappus, IV, and H. Hayes. Determining credibility from social network structure. In *ASONAM'13*.
- [4] A. Z. Broder, S. C. Glassman, M. S. Manasse, and G. Zweig. Syntactic clustering of the web. In *Computer Networks and ISDN Systems'97*.
- [5] R. S. Burt. Structural holes: The social structure of competition. In *Explorations in economic sociology'09*.
- [6] J. M. Campanario. Empirical study of journal impact factors obtained using the classical two-year citation window versus a five-year citation window. In *Scientometrics'11*.
- [7] Y. Cao, C.-Y. Chan, J. Li, and K.-L. Tan. Optimization of analytic window functions. In *VLDB'12*.
- [8] C. Chen, X. Yan, F. Zhu, J. Han, and P. S. Yu. Graph olap: Towards online analytical processing on graphs. In *ICDM'08*.
- [9] J. Cheng, Z. Shang, H. Cheng, H. Wang, and J. X. Yu. K-reach: who is in your small world. In *VLDB'12*.
- [10] A. Czumaj, M. Kowaluk, and A. Lingas. Faster algorithms for finding lowest common ancestors in directed acyclic graphs. *Theoretical Computer Science*, 2007.
- [11] L. Dai, J.-d. Luo, X. Fu, and Z. Li. Predicting offline behaviors from online features: An ego-centric dynamical

network approach.

- [12] C. W. Holsapple and W. Luo. A citation analysis of influences on collaborative computing research. In *Comput. Supported Coop. Work'03*.
- [13] I. M. Keseler, J. Collado-Vides, S. Gama-Castro, J. Ingraham, S. Paley, I. T. Paulsen, M. Peralta-Gil, and P. D. Karp. Ecocyc: a comprehensive database resource for escherichia coli. In *Nucleic acids research'05*.
- [14] H. H. Ma, S. Gustafson, A. Moitra, and D. Bracewell. Ego-centric network sampling in viral marketing applications. In *CSE'09*.
- [15] N. Ma, J. Guan, and Y. Zhao. Bringing pagerank to the citation analysis. In *Inf. Process. Manage.'08*.
- [16] J. Mondal and A. Deshpande. Eagr: Supporting continuous ego-centric aggregate queries over large dynamic graphs. In *SIGMOD'14*.
- [17] Y. Tian, R. A. Hankins, and J. M. Patel. Efficient aggregation for graph summarization. In *SIGMOD'08*.
- [18] V. Vassilevska and A. Pinar. Finding nonoverlapping dense blocks of a sparse matrix. In *Lawrence Berkeley National Laboratory'04*.
- [19] Z. Wang, Q. Fan, H. Wang, K.-l. Tan, D. Agrawal, and A. El Abbadi. Pagrol: Parallel graph olap over large-scale attributed graphs. In *ICDE'14*.
- [20] X. Yan, B. He, F. Zhu, and J. Han. Top-k aggregation queries over large networks. In *ICDE'10*.
- [21] H. Yildirim, V. Chaoji, and M. J. Zaki. Dagger: A scalable index for reachability queries in large dynamic graphs. 2013.
- [22] J. X. Yu and J. Cheng. Graph reachability queries: A survey. In *Managing and Mining Graph Data*. 2010.
- [23] F. Zemke. What's new in sql: 2011. In *ACM SIGMOD Record'12*.
- [24] P. Zhao, X. Li, D. Xin, and J. Han. Graph cube: on warehousing and olap multidimensional networks. In *SIGMOD'11*.